arXiv:0711.0185v1 [math.NT] 1 Nov 2007

THE TRUE COMPLEXITY OF A SYSTEM OF LINEAR EQUATIONS

W.T. GOWERS AND J. WOLF

Abstract. In this paper we look for conditions that are sufficient to guarantee that a subset A of a finite Abelian group G contains the "expected" number of linear configurations of a given type. The simplest non-trivial result of this kind is the well-known fact that if G has odd order, A has density α and all Fourier coefficients of the characteristic function of A are significantly smaller than α (except the one at zero, which equals α), then A contains approximately $\alpha^3|G|^2$ triples of the form $(a, a+d, a+2d)$. This is "expected" in the sense that a random set A of density α has approximately $\alpha^3|G|^2$ such triples with very high probability.

More generally, it was shown in [Gow01] (in the case $G = \mathbb{Z}_N$ for N prime, but the proof generalizes) that a set A of density α has about $\alpha^k|G|^2$ arithmetic progressions of length k if the characteristic function of A is almost as small as it can be, given its density, in a norm that is now called the U^{k-1}-norm. Green and Tao [GT06] have found the most general statement that follows from the technique used to prove this result, introducing a notion that they call the *complexity* of a system of linear forms. They prove that if A has almost minimal U^{k+1}-norm then it has the expected number of linear configurations of a given type, provided that the associated complexity is at most k. The main result of this paper is that the converse is not true: in particular there are certain systems of complexity 2 that are controlled by the U^2-norm, whereas the result of Green and Tao requires the stronger hypothesis of U^3-control.

We say that a system of m linear forms $L_1, \dots, L_m$ in d variables has *true complexity* k if k is the smallest positive integer such that, for any set A of density α and almost minimal U^{k+1}-norm, the number of d-tuples $(x_1, \dots, x_d)$ such that $L_i(x_1, \dots, x_d) \in A$ for every i is approximately $\alpha^m|G|^d$. We conjecture that the true complexity k is the smallest positive integer s for which the functions $L_1^{s+1}, \dots, L_m^{s+1}$ are linearly independent. Using the "quadratic Fourier analysis" of Green and Tao we prove this conjecture in the case where the complexity of the system (in Green and Tao's sense) is 2, $s = 1$ and G is the group $\mathbb{F}_p^n$ for some fixed odd prime p.

A closely related result in ergodic theory was recently proved independently by Leibman [Lei07]. We end the paper by discussing the connections between his result and ours.

Date: 1st November 2007.

1. Introduction

Let A be a subset of a finite Abelian group G and let $\alpha = |A|/|G|$ be the *density* of A. We say that A is *uniform* if it has one of several equivalent properties, each of which says in its own way that A "behaves like a random set". For example, writing A for the characteristic function of the set A, we can define the convolution $A * A$ by the formula

$$A * A(x) = \mathbb{E}_{y+z=x} A(y)A(z),$$

where the expectation is with respect to the uniform distribution over all pairs $(y, z) \in G^2$ such that $y + z = x$; one of the properties in question is that the variance of $A * A$ should be small. If this is the case and G has odd order, then it is easy to show that A contains approximately $\alpha^3|G|^2$ triples of the form $(x, x + d, x + 2d)$. Indeed, these triples are the solutions (x, y, z) of the equation $x + z = 2y$, and

$$\mathbb{E}_{x+z=2y} A(x)A(y)A(z) = \mathbb{E}_y A * A(2y)A(y).$$

The mean of the function $A * A$ is α^2, so if the variance is sufficiently small, then the right-hand side is approximately $\alpha^2\mathbb{E}_y A(y) = \alpha^3$. This is a probabilistic way of saying that the number of solutions of $x + z = 2y$ inside A is approximately $\alpha^3|G|^2$, which is what we would expect if A was a random set with elements chosen independently with probability α.

An easy generalization of the above argument shows that, given any linear equation in G of the form

$$c_1x_1 + c_2x_2 + \cdots + c_mx_m = 0,$$

for suitable fixed coefficients $c_1, c_2, ..., c_m$, the number of solutions in A is approximately $\alpha^m|G|^{m-1}$. Roughly speaking, you can choose $x_3, \ldots, x_m$ in A however you like, and if A is sufficiently uniform then the number of ways of choosing x_1 and x_2 to lie in A and satisfy the equation will almost always be roughly $\alpha^2|G|$. By "suitable" we mean that there are certain divisibility problems that must be avoided. For example, if G is the group $\mathbb{F}_2^n$, $x + z = 2y$ and x belongs to A, then z belongs to A for the trivial reason that it equals x. Throughout this paper we shall consider groups of the form $\mathbb{F}_p^n$ for some prime p and assume that p is large enough for such problems not to arise.

When $k \geq 4$, uniformity of a set A does not guarantee that A contains approximately $\alpha^k|G|^2$ arithmetic progressions of length k. For instance, there are examples of uniform subsets of $\mathbb{Z}_N$ that contain significantly more, or even significantly fewer than, the expected

number of four-term progressions [Gow06]. It was established in [Gow98] that the appropriate measure for dealing with progressions of length 4 is a property known as *quadratic uniformity*: sets which are sufficiently quadratically uniform contain roughly the correct number of four-term progressions. We shall give precise definitions of higher-degree uniformity in the next section, but for now let us simply state the result, proved in [Gow01] in the case $G = \mathbb{Z}_N$, that if A is uniform of degree $k-2$, then A contains approximately $\alpha^k|G|^2$ arithmetic progressions of length k. Moreover, if A is uniform of degree j for some $j < k-2$, then it does *not* follow that A must contain approximately $\alpha^k|G|^2$ arithmetic progressions of length k.

The discrepancy between k and $k-2$ seems slightly unnatural until one reformulates the statement in terms of solutions of equations. We can define an arithmetic progression of length k either as a k-tuple of the form $(x, x+d, \dots, x+(k-1)d)$ or as a solution $(x_1, x_2, \dots, x_k)$ to the system of $k-2$ equations $x_i - 2x_{i+1} + x_{i+2} = 0$, $i = 1, 2, \dots, k-2$. In all the examples we have so far discussed, we need uniformity of degree precisely k in order to guarantee approximately the expected number of solutions of a system of k equations. It is tempting to ask whether this is true in general.

However, a moment's reflection shows that it is not. For example, the system of equations $x_1 - 2x_2 + x_3 = 0$, $x_4 - 2x_5 + x_6 = 0$ has about $\alpha^6|G|^4$ solutions in a uniform set, since the two equations are completely independent. This shows that a sensible conjecture must take account of how the equations interact with each other.

A more interesting example is the system that consists of the $\binom{m}{3}$ equations $x_{ij}+x_{jk} = x_{ik}$ in the $\binom{m}{2}$ unknowns x_{ij}, $1 \leq i < j \leq m$. These equations are not all independent, but one can of course choose an independent subsystem of them. It is not hard to see that there is a bijection between solutions of this system of equations where every x_{ij} belongs to A and m-tuples $(x_1, \dots, x_m)$ such that $x_j - x_i \in A$ whenever $i < j$. Now one can form a bipartite graph with two vertex sets equal to G by joining x to y if and only if $y - x \in A$. It is well-known that if A is uniform, then this bipartite graph is quasirandom. The statement that every $x_j - x_i$ belongs to A can be reformulated to say that $(x_1, \dots, x_m)$ form a clique in an m-partite graph that is built out of quasirandom pieces derived from A. A "counting lemma" from the theory of quasirandom graphs then implies easily that the number of such cliques is approximately $\alpha^{\binom{m}{2}}|G|^m$. So uniformity of degree 1 is sufficient to guarantee that there are about the expected number of solutions to this fairly complicated system of equations.

In their recent work on configurations in the primes, Green and Tao [GT06] analysed the arguments used to prove the above results, which are fairly simple and based on repeated use of the Cauchy-Schwarz inequality. They isolated the property that a system of equations, or equivalently a system of linear forms, must have in order for degree-k uniformity to be sufficient for these arguments to work, and called this property *complexity*. Since in this paper we shall have more than one notion of complexity, we shall sometimes call their notion *Cauchy-Schwarz complexity*, or *CS-complexity* for short.

Definition 1.1. *Let $\mathcal{L} = (L_1, ..., L_m)$ be a system of m linear forms in d variables. For $1 \leq i \leq m$ and $s \geq 0$, we say that $\mathcal{L}$ is* s-complex *at i if one can partition the $m-1$ forms $\{L_j : j \neq i\}$ into $s+1$ classes such that L_i does not lie in the linear span of any of these classes. The* Cauchy-Schwarz complexity *(or CS-complexity) of $\mathcal{L}$ is defined to be the least s for which the system is s-complex at i for all $1 \leq i \leq m$, or ∞ if no such s exists.*

To get a feel for this definition, let us calculate the complexity of the system $\mathcal{L}$ of k linear forms $x, x+y, \dots, x+(k-1)y$. Any two distinct forms $x+iy$ and $x+jy$ in $\mathcal{L}$ contain x and y in their linear span. Therefore, whichever form L we take from $\mathcal{L}$, if we wish to partition the others into classes that do not contain L in their linear span, then we must take these classes to be singletons. Since we are partitioning $k-1$ forms, this tells us that the minimal s is $k-2$. So $\mathcal{L}$ has complexity $k-2$.

Next, let us briefly look at the system $\mathcal{L}$ of $\binom{m}{2}$ forms $x_i - x_j$ $(1 \leq i < j \leq m)$ that we discussed above. If L is the form $x_i - x_j$ then no other form $L' \in \mathcal{L}$ involves both x_i and x_j, so we can partition $\mathcal{L} \setminus \{L\}$ into the forms that involve x_i (which therefore do not involve x_j) and the forms that do not involve x_i. Since neither class includes L in its linear span, the complexity of $\mathcal{L}$ is at most 1. When $m \geq 3$ it can also be shown to be at least 1.

It follows from Green and Tao's result that if A is sufficiently uniform and $\mathcal{L} = (L_1, ..., L_m)$ has complexity at most 1, then A contains approximately the expected number of m-tuples of the form $(L_1(x_1, \dots, x_d), \dots, L_m(x_1, \dots x_d))$. (If the forms are defined over $\mathbb{Z}^d$, then this number is $\alpha^m |G|^d$.)

Notice that this statement adequately explains all the cases we have so far looked at in which uniformity implies the correct number of solutions. It is thus quite natural to conjecture that Green and Tao's result is tight. That is, one might guess that if the complexity $\mathcal{L}$ is greater than 1 then there exist sets A that do not have the correct number of images of $\mathcal{L}$.

But is this correct? Let us look at what is known in the other direction, by discussing briefly the simplest example that shows that uniform sets in $\mathbb{Z}_N$ do not have to contain

the correct number of arithmetic progressions of length 4. (Here we are taking N to be some large prime.) Roughly speaking, one takes A to be the set of all x such that x^2 mod N is small. Then one makes use of the identity

$$x^2 - 3(x+d)^2 + 3(x+2d)^2 - (x+3d)^2 = 0$$

to prove that if x, $x+d$ and $x+2d$ all lie in A, then $x+3d$ is rather likely to lie in A as well, because $(x+3d)^2$ is a small linear combination of small elements of $\mathbb{Z}_N$. This means that A has "too many" progressions of length 4. (Later, we shall generalize this example and make it more precise.)

The above argument uses the fact that the squares of the linear forms x, $x+d$, $x+2d$ and $x+3d$ are linearly dependent. Later, we shall show that if $\mathcal{L}$ is *any* system of linear forms whose squares are linearly dependent, then essentially the same example works for $\mathcal{L}$. This gives us a sort of "upper bound" for the set of systems $\mathcal{L}$ that have approximately the right number of images in any uniform set: because of the above example, we know that the squares of the forms in any such system $\mathcal{L}$ must be linearly independent.

And now we arrive at the observation that motivated this paper: the "upper bound" just described does not coincide with the "lower bound" of Green and Tao. That is, there are systems of linear forms of complexity greater than 1 with squares that are linearly independent. One of the simplest examples is the system $(x, y, z, x+y+z, x+2y-z, x+2z-y)$. Another, which is translation-invariant (in the sense that if you add a constant to everything in the configuration, you obtain another configuration of the same type), is $(x, x+y, x+z, x+y+z, x+y-z, x+z-y)$. A third and rather natural example that is also translation-invariant is the configuration

$$(x, x+a, x+b, x+c, x+a+b, x+a+c, x+b+c),$$

which can be thought of as a cube minus a point. All these examples have complexity 2, but it is not hard to produce examples with arbitrarily high complexity.

In the light of such examples, we are faced with an obvious question: which systems of linear forms have roughly the expected number of images in any sufficiently uniform set? We conjecture that the correct answer is given by the "upper bound"—that is, that square independence is not just necessary but also sufficient. When the group G is $\mathbb{F}_p^n$ for a fixed prime p, we prove this conjecture for systems of complexity 2. This includes the two examples above, and shows that having Cauchy-Schwarz complexity at most 1 is not a necessary condition, even if it is a natural sufficient one.

However, the proof is much deeper for systems of complexity 2. Although the statement of our result is completely linear, we use "quadratic Fourier analysis", recently developed by Green and Tao [GT05a], to prove it, and it seems that we are forced to do so. Thus, it appears that Cauchy-Schwarz complexity captures the systems for which an easy argument exists, while square independence captures the systems for which the result is true.

Very recently, and independently, Leibman [Lei07] described a similar phenomenon in the ergodic-theoretic context. In the final section of the paper we shall briefly outline how his results relate to ours.

So far, we have concentrated on uniform sets. However, in the next section we shall define higher-degree uniformity and formulate a more complete conjecture, which generalizes the above discussion in a straightforward way. Green and Tao proved that a system of Cauchy-Schwarz complexity k has approximately the correct number of images in a set A if A is sufficiently uniform of degree $k+1$. Once again, it seems that this is not the whole story, and that the following stronger statement should be true: a linear system $\mathcal{L} = (L_1, \dots, L_m)$ has the right number of images in any set A that is sufficiently uniform of degree k if and only if the functions L_i^{k+1} are linearly independent. The reason we have not proved this is that the natural generalization of our existing argument would have to use an as yet undeveloped general "polynomial Fourier analysis", which is known only in the quadratic case. However, it is easy to see how our arguments would generalize if such techniques were available, which is compelling evidence that our conjecture (which we will state formally in a moment) is true.

2. Uniformity norms and true complexity

As promised, let us now give a precise definition of higher-degree uniformity. We begin by defining a sequence of norms, known as *uniformity norms*.

Definition 2.1. *Let G be a finite Abelian group. For any positive integer $k \geq 2$ and any function $f : G \to \mathbb{C}$, define the U^k-*norm *by the formula*

$$\|f\|_{U^k}^{2^k} := \mathbb{E}_{x,h_1,\dots,h_k \in G} \prod_{\boldsymbol{\omega} \in \{0,1\}^k} C^{|\boldsymbol{\omega}|} f(x + \boldsymbol{\omega} \cdot \mathbf{h}),$$

where $\boldsymbol{\omega} \cdot \mathbf{h}$ is shorthand for $\sum_i \omega_i h_i$, and $C^{|\boldsymbol{\omega}|} f = f$ if $\sum_i \omega_i$ is even and $\overline{f}$ otherwise.

These norms were first defined in [Gow01] (in the case where G is the group $\mathbb{Z}_N$). Of particular interest in this paper will be the U^2-norm and the U^3-norm. The former can be described in many different ways. The definition above expresses it as the fourth root of

the average of

$$f(x)\overline{f(x+h)f(x+h')}f(x+h+h')$$

over all triples (x,h,h'). It is not hard to show that this average is equal to $\|f*f\|_2^2$, and also to $\|\hat{f}\|_4^4$. (These identities depend on appropriate normalizations—we follow the most commonly used convention of taking averages in physical space and sums in frequency space.)

We shall call a function f *c-uniform* if $\|f\|_{U^2}\leq c$ and *c-quadratically uniform* if $\|f\|_{U^3}\leq c$. We shall often speak more loosely and describe a function as uniform if it is c-uniform for some small c, and similarly for higher-degree uniformity. We remark here that if $j\leq k$ then $\|f\|_{U^j}\leq\|f\|_{U^k}$, so c-uniformity of degree k implies c-uniformity of all lower degrees.

If A is a subset of an Abelian group G and the density of A is α, then we say that A is uniform of degree k if it is close in the U^k-norm to the constant function α. More precisely, we define the *balanced function* $f(x)=A(x)-\alpha$ and say that A is *c-uniform of degree k* if $\|f\|_{U^k}\leq c$.

The following theorem is essentially Theorem 3.2 in [Gow01]. (More precisely, in that paper the theorem was proved for the group $\mathbb{Z}_N$, but the proof is the same.)

Theorem 2.2. *Let $k\geq 2$ and let G be a finite Abelian group such that there are no non-trivial solutions to the equation $jx=0$ for any $1\leq j<k$. Let $c>0$ and let $f_1,f_2,\dots,f_k$ be functions from G to $\mathbb{C}$ such that $\|f_i\|_\infty\leq 1$ for every i. Then*

$$\Big|\mathbb{E}_{x,y\in G}f_1(x)f_2(x+y)\dots f_k(x+(k-1)y)\Big|\leq\|f_k\|_{U^{k-1}}.$$

It follows easily from this result that if A is a set of density α and A is c-uniform for sufficiently small c, then A contains approximately $\alpha^k|G|^2$ arithmetic progressions of length k. Very briefly, the reason for this is that we are trying to show that the average

$$\mathbb{E}_{x,y}A(x)A(x+y)\dots A(x+(k-1)y)$$

is close to α^k. Now this average is equal to

$$\mathbb{E}_{x,y}A(x)A(x+y)\dots f(x+(k-1)y)+\alpha\mathbb{E}_{x,y}A(x)A(x+y)\dots A(x+(k-2)y).$$

The first of these terms is at most c, by Theorem 2.2, and the second can be handled inductively. The bound we obtain in this way is $c(1+\alpha+\dots+\alpha^{k-1})\leq kc$.

We can now state formally Green and Tao's generalization in terms of CS-complexity in the case where G is the group $\mathbb{Z}_N$, which is implicit in [GT06].

Theorem 2.3. *Let N be a prime, let $f_1, \dots, f_m$ be functions from $\mathbb{Z}_N$ to $[-1,1]$, and let $\mathcal{L}$ be a linear system of CS-complexity k consisting of m forms in d variables. Then, provided $N \geq k$,*

$$\left|\mathbb{E}_{x_1,\dots,x_d \in \mathbb{Z}_N} \prod_{i=1}^{m} f(L_i(x_1,\dots,x_d))\right| \leq \min_i \|f_i\|_{U^{k+1}}.$$

Just as in the case of arithmetic progressions, it follows easily that if A is a subset of G of density α, then the probability, given a random element $(x_1,\dots,x_d) \in G^d$, that all the m images $L_i(x_1,\dots,x_d)$ lie in A is approximately α^m. (The inductive argument depends on the obvious fact that if $\mathcal{L}$ has complexity at most k then so does any subsystem of $\mathcal{L}$.)

Green and Tao proved the above theorem because they were investigating which linear configurations can be found in the primes. For that purpose, they in fact needed a more sophisticated "relative" version of the statement. Since the proof of the version we need here is simpler (partly because we are discussing systems of complexity at most 2, but much more because we do not need a relative version), we give it for the convenience of the reader. This is another result where the proof is essentially the same for all Abelian groups, give or take questions of small torsion. Since we need it in the case $G = \mathbb{F}_p^n$, we shall just prove it for this group. The reader should bear in mind that for this group, one should understand linear independence of a system of forms as independence over $\mathbb{F}_p$ when one is defining complexity (and also square-independence).

The first step of Green and Tao's proof was to put an arbitrary linear system into a convenient form for proofs. Given a linear form L in d variables $x_1, \dots, x_d$, let us define the *support* of L to be the set of j such that L depends on x_j. That is, if $L(x_1,\dots,x_d) = \lambda_1 x_1 + \cdots + \lambda_d x_d$ then the support of L is $\{i : \lambda_i \neq 0\}$. Let $\mathcal{L} = (L_1,\dots,L_m)$ be a system of linear forms and let the support of L_i be σ_i for each i. Then $\mathcal{L}$ is said to be in *s-normal form* if it is possible to find subsets $\tau_i \subset \sigma_i$ for each i with the following two properties.

(i) Each τ_i has cardinality at most $s+1$.

(ii) If $i \neq j$ then τ_i is not a subset of σ_j.

If a linear system $\mathcal{L}$ is in s-normal form, then it has complexity at most s. Indeed, if τ_i has r elements $\{i_1,\dots,i_r\}$, then one can partition the remaining forms into r sets $\mathcal{L}_1,\dots,\mathcal{L}_r$ in such a way that no form in $\mathcal{L}_h$ uses the variable x_{i_h}. Since L_i *does* use the variable x_{i_h} it is not in the linear span of $\mathcal{L}_h$.

The converse of this statement is false, but Green and Tao prove that every linear system of complexity s can be "extended" to one that is in s-normal form. This part of the proof

is the same in both contexts, so we do not reproduce it. All we need to know here is that if we prove Theorem 2.3 for systems in normal form then we have it for general systems.

Just to illustrate this, consider the obvious system associated with arithmetic progressions of length 4, namely $(x, x+y, x+2y, x+3y)$. This is not in 2-normal form, because the support of the first form is contained in the supports of the other three. However, the system $(-3x-2y-z, -2x-y+w, -x+z+2w, y+2z+3w)$ *is* in 2-normal form (since the supports have size 3 and are distinct) and its images are also uniformly distributed over all arithmetic progressions of length 4 (if we include degenerate ones).

Now let us prove Theorem 2.3 when $k = 2$. Without loss of generality we may assume that $\mathcal{L}$ is in 2-normal form at 1, and that it is the only form using all three variables $x_1 = x, x_2 = y$ and $x_3 = z$. We use the shorthand $h(x, y, z) = f(L_1(x_1, x_2, ..., x_d))$, and denote by $b(x, y)$ any general bounded function in two variables x and y. It is then possible to rewrite

$$\mathbb{E}_{x_1,...,x_d \in \mathbb{F}_p^n} \prod_{i=1}^{m} f(L_i(x_1, ..., x_d))$$

as

$$\mathbb{E}_{x_4,x_5,...,x_d} \mathbb{E}_{x,y,z} h(x, y, z)b(x, y)b(y, z)b(x, z).$$

Here, the functions h and b depend on the variables $x_4, \dots, x_d$ but we are suppressing this dependence in the notation.

Estimating the expectation over (x, y, z) is a well-known argument from the theory of quasirandom hypergraphs. (See for instance Theorem 4.1 in [Gow04].) First, we apply Cauchy-Schwarz and use the boundedness of b to obtain an upper bound of

$$(\mathbb{E}_{x,y}(\mathbb{E}_z h(x, y, z)b(x, z)b(y, z))^2)^{1/2}.$$

Expanding out the square and rearranging yields

$$(\mathbb{E}_{y,z,z'} b(y, z)b(y, z')\mathbb{E}_x h(x, y, z)h(x, y, z')b(x, z)b(x, z'))^{1/2},$$

and by a second application of Cauchy-Schwarz we obtain an upper bound of

$$(\mathbb{E}_{y,z,z'}(\mathbb{E}_x h(x, y, z)h(x, y, z')b(x, z)b(x, z'))^2)^{1/4}.$$

A second round of interchanging summation followed by a third application of Cauchy-Schwarz gives us an upper bound of

$$(\mathbb{E}_{x,x',z,z'}(\mathbb{E}_y h(x, y, z)h(x, y, z')h(x', y, z)h(x', y, z'))^2)^{1/8}.$$

This expression equals the "octahedral norm" of the function $h(x,y,z)$—a hypergraph analogue of the U^3-norm. Because for fixed $x_4,\dots,x_d$, h depends only on the linear expression $L_1(x,y,z)$, a simple change of variables can be used to show that it is in fact equal to $\|f\|_{U^3}$.

Now all that remains is to take the expectation over the remaining variables and the proof is complete. It is also not hard to generalize to arbitrary k, but this we leave as an exercise to the reader.

Now, as we stated earlier, Theorem 2.3 does not settle the question of which systems are controlled by which degrees of uniformity. Accordingly, we make the following definition.

Definition 2.4. *Let $\mathcal{L}$ be a system of m distinct linear forms $L_1, L_2, \dots, L_m$ in d variables. The* true complexity *of $\mathcal{L}$ is the smallest k with the following property. For every $\epsilon > 0$ there exists $\delta > 0$ such that if G is any finite Abelian group and $f : G \to \mathbb{C}$ is any function with $\|f\|_\infty \leq 1$ and $\|f\|_{U^{k+1}} \leq \delta$, then*

$$\left|\mathbb{E}_{x_1,\dots,x_d \in G} \prod_{i=1}^{m} f(L_i(x_1,\dots,x_d))\right| \leq \epsilon.$$

The main conjecture of this paper is now simple to state precisely.

Conjecture 2.5. *The true complexity of a system of linear forms $\mathcal{L} = (L_1,\dots,L_m)$ is equal to the smallest k such that the functions L_i^{k+1} are linearly independent.*

In the next section, we shall prove this conjecture in the simplest case that is not covered by the result of Green and Tao, namely the case when $k = 1$ and $\mathcal{L}$ has CS-complexity 2. All other cases would require a more advanced form of polynomial Fourier analysis than the quadratic Fourier analysis that is so far known, but we shall explain why it will almost certainly be possible to generalize our argument once such a theory is developed.

3. True complexity for vector spaces over finite fields

We shall now follow the course that is strongly advocated by Green [Gre05a] and restrict attention to the case where G is the group $\mathbb{F}_p^n$, where p is a fixed prime and n tends to infinity. The reason for this is that it makes many arguments technically simpler than they are for groups with large torsion such as $\mathbb{Z}_N$. In particular, one can avoid the technicalities associated with Bohr sets. These arguments can then almost always be converted into more complicated arguments for $\mathbb{Z}_N$. (In a forthcoming paper, we give a different proof for the case $\mathbb{F}_p^n$ and carry out the conversion process. That proof is harder than the proof here but gives significantly better bounds and is easier to convert.)

We begin this section with the easier half of our argument, showing that if $\mathcal{L}$ is a system of linear forms $(L_1, \dots, L_m)$ and if there is a linear dependence between the squares of these forms, then the true complexity of $\mathcal{L}$ is greater than 1. This part can be proved almost as easily for $\mathbb{Z}_N$, but we shall not do so here.

3.1. **Square-independence is necessary.** Let us start by briefly clarifying what we mean by square-independence of a linear system $\mathcal{L} = (L_1, \dots, L_m)$. When the group G is $\mathbb{Z}_N$, then all we mean is that the functions L_i^2 are linearly independent, but when it is $\mathbb{F}_p^n$, then this definition does not make sense any more. Instead, we ask for the quadratic forms $L_i^T L_i$ to be linearly independent. If $L_i(x_1, \dots, x_d) = \sum_r \gamma_r^{(i)} x_r$, then $L_i^T L_i(x_1, \dots, x_d) = \sum_r \sum_s \gamma_r^{(i)} \gamma_s^{(i)} x_r x_s$. Therefore, what we are interested in is linear independence of the matrices $\Gamma_{rs}^{(i)} = \gamma_r^{(i)} \gamma_s^{(i)}$ over $\mathbb{F}_p$. (Note that in the case of $\mathbb{Z}_N$, this is equivalent to independence of the functions L_i^2.)

Theorem 3.1. *Let $\mathcal{L} = (L_1, \dots, L_m)$ be a system of linear forms in d variables and suppose that the quadratic forms $L_i^T L_i$ are linearly dependent over $\mathbb{F}_p$. Then there exists $\epsilon > 0$ such that for every $\delta > 0$ there exists n and a set $A \subset \mathbb{F}_p^n$ with the following two properties.*

(i) A is δ-uniform of degree 1.

(ii) If $\mathbf{x} = (x_1, \dots, x_d)$ is chosen randomly from $(\mathbb{F}_p^n)^d$, then the probability that $L_i(\mathbf{x})$ is in A for every i is at least $\alpha^m + \epsilon$, where α is the density of A.

In other words, the true complexity of $\mathcal{L}$ is at least 2.

For the proof we require the following standard lemma, which says that certain Gauss sums are small. A proof can be found in [Gre05b], for example.

Lemma 3.2. *Suppose that $q : \mathbb{F}_p^n \to \mathbb{F}_p$ is a quadratic form of rank r. That is, suppose that $q(x) = x^T M x + b^T x$ for some matrix M of rank r and some vector $b \in \mathbb{F}_p^n$. Then*

$$|\mathbb{E}_{x \in \mathbb{F}_p^n} \omega^{q(x)}| \leq p^{-r/2},$$

with equality if $b = 0$. In particular,

$$|\mathbb{E}_{x \in \mathbb{F}_p^n} \omega^{\eta x^T x}| \leq p^{-n/2}$$

for any non-zero $\eta \in \mathbb{F}_p$.

Proof of Theorem 3.1. Let A be the set $\{x \in \mathbb{F}_p^n : x^T x = 0\}$. Then the characteristic function of A can be written as

$$A(x) = \mathbb{E}_u \omega^{u x^T x},$$

where $\omega = \exp(2\pi i/p)$ and the expectation is taken over $\mathbb{F}_p$. Let us now take any square-*independent* system $\mathcal{L} = (L_1, \dots, L_m)$ of linear forms in $\mathbf{x} = (x_1, \dots, x_d)$ and estimate the expectation $\mathbb{E}_{\mathbf{x}} \prod_i A(L_i(\mathbf{x}))$.

Using the formula for $A(x)$, we can rewrite this expectation as

$$\mathbb{E}_{\mathbf{x}\in(\mathbb{F}_p^n)^d} \mathbb{E}_{u_1,\dots,u_m\in\mathbb{F}_p} \omega^{\sum_i u_i L_i(\mathbf{x})^T L_i(\mathbf{x})}.$$

We can break this up into p^m expectations over $\mathbf{x}$, one for each choice of $u_1, \dots, u_m$.

If the u_i are all zero, then the expectation over $\mathbf{x}$ is just the expectation of the constant function 1, so it is 1. Otherwise, since the quadratic forms $L_i^T L_i$ are linearly independent, the sum $\sum_i u_i L_i(\mathbf{x})^T L_i(\mathbf{x})$ is a non-zero quadratic form $q(x) = \sum_{i,j} \gamma_{ij} x_i^T x_j$.

Without loss of generality, there exists j such that $\gamma_{1j} \neq 0$. If in addition $\gamma_{11} = 0$, then for every choice of $x_2, \dots, x_d$ we can write $q(x)$ in the form $r^T x_1 + z$, where $r = \sum_j \gamma_{1j} x_j$ and z depends on $x_2, \dots, x_d$ only. This is a non-constant linear function of x_1 except when $\sum_j \gamma_{1j} x_j = 0$. Since not every γ_{1j} is zero, this happens with probability p^{-n}. Therefore, $|\mathbb{E}_{\mathbf{x}} \omega^{q(\mathbf{x})}| \leq p^{-n}$ in this case. If $\gamma_{11} \neq 0$, then this same function has the form $\gamma_{11} x_1^T x_1 + r^T x_1$ for some element $r \in \mathbb{F}_p^n$ (which depends on $x_2, \dots, x_d$). In this case, Lemma 3.2 implies that the expectation is at most $p^{-n/2}$.

Since the probability that $u_1 = \cdots = u_m = 0$ is p^{-m}, this shows that

$$\left| \mathbb{E}_{\mathbf{x}\in(\mathbb{F}_p^n)^d} \prod_i A(L_i(\mathbf{x})) - p^{-m} \right| \leq p^{-n/2}.$$

Applying this result in the case where $\mathcal{L}$ consists of the single form x, we see that the density of A differs from p^{-1} by at most $p^{-n/2}$. Therefore, we have shown that for this particular set A, square-independence of $\mathcal{L}$ guarantees approximately the "correct" probability that every $L_i(\mathbf{x})$ lies in A.

This may seem like the opposite of what we were trying to prove, but in fact we have almost finished, for the following simple reason. If we now take $\mathcal{L}$ to be an arbitrary system $(L_1, \dots, L_m)$ of linear forms, then we can choose from it a maximal square-independent subsystem. Without loss of generality this subsystem is $(L_1, \dots, L_l)$. Then all the quadratic forms $L_i^T L_i$ with $i > l$ are linear combinations of $L_1^T L_1, \dots, L_l^T L_l$, so a sufficient condition for every $L_i^T L_i(\mathbf{x})$ to be zero is that it is zero for every $i \leq l$. But this we know happens with probability approximately p^{-l} by what we have just proved. Therefore, if $\mathcal{L}$ is not square-independent, then A^m contains "too many" m-tuples of the form $(L_1(\mathbf{x}), \dots, L_m(\mathbf{x}))$. □

3.2. **A review of quadratic Fourier analysis.** We shall now turn our attention to the main result of this paper, which states that if $\mathcal{L}$ has CS-complexity at most 2 and is

square-independent, then the true complexity of $\mathcal{L}$ is at most 1. We begin with a quick review of quadratic Fourier analysis for functions defined on $\mathbb{F}_p^n$. Our aim in this review is to give precise statements of the results that we use in our proof. The reader who is prepared to use quadratic Fourier analysis as a black box should then find that this paper is self-contained.

So far in our discussion of uniformity, we have made no mention of Fourier analysis at all. However, at least for the U^2-norm, there is a close connection. Let f be a complex-valued function defined on a finite Abelian group G. If χ is a character on G, the Fourier coefficient $\hat{f}(\chi)$ is defined to be $\mathbb{E}_x f(x)\chi(x)$. The resulting Fourier transform satisfies the convolution identity $\widehat{f * g} = \hat{f}\hat{g}$, Parseval's identity $\|\hat{f}\|_2 = \|f\|_2$ and the inversion formula $f(x) = \sum_\chi \hat{f}(\chi)\chi(-x)$. (The second and third identities depend on the correct choice of normalization: $\|f\|_2^2$ is defined to be $\mathbb{E}_x|f(x)|^2$, whereas $\|\hat{f}\|_2^2$ is defined to be $\sum_\chi |\hat{f}(\chi)|^2$. That is, as mentioned earlier, we take averages in G and sums in $\hat{G}$.) It follows that $\|f\|_{U^2}^4 = \|\hat{f}\|_4^4$, since both are equal to $\|f * f\|_2^2$.

It is often useful to split a function f up into a "structured" part and a uniform part. One way of doing this is to let K be the set of all characters χ for which $|\hat{f}(\chi)|$ is larger than some δ and to write $f = f_1 + f_2$, where $f_1 = \sum_{\chi \in K} \hat{f}(\chi)\chi(-x)$ and $f_2 = \sum_{\chi \notin K} \hat{f}(\chi)\chi(-x)$. If $\|f\|_\infty \leq 1$, (as it is in many applications), then Parseval's identity implies that $|K| \leq \delta^{-2}$, and can also be used to show that $\|f_2\|_{U^2} \leq \delta^{1/2}$. That is, K is not too large, and f_2 is $\delta^{1/2}$-uniform.

When G is the group $\mathbb{F}_p^n$, the characters all have the form $x \mapsto \omega^{r^T x}$. Notice that this character is constant on all sets of the form $\{x : r^T x = u\}$, and that these sets partition $\mathbb{F}_p^n$ into p affine subspaces of codimension 1. Therefore, one can partition $\mathbb{F}_p^n$ into at most $p^{|K|}$ affine subspaces of codimension $|K|$ such that f_1 is constant on each of them. This is the sense in which f_1 is "highly structured".

The basic aim of quadratic Fourier analysis is to carry out a similar decomposition for the U^3-norm. That is, given a function f, we would like to write f as a sum $f_1 + f_2$, where f_1 is "structured" and f_2 is *quadratically* uniform. Now this is a stronger (in fact, much stronger) property to demand of f_2, so we are forced to accept a weaker notion of structure for f_1.

Obtaining any sort of structure at all is significantly harder than it is for the U^2-norm, and results in this direction are much more recent. The first steps were taken in [Gow98] and [Gow01] for the group $\mathbb{Z}_N$ in order to give an analytic proof of Szemerédi's theorem. The structure of that proof was as follows: Theorem 2.2 (of the present paper) can be

used to show that if a set A is sufficiently uniform of degree $k-2$, then it must contain an arithmetic progression of length k. Then an argument that is fairly easy when $k=3$ but much harder when $k \geq 4$ can be used to show that if A is *not* c-uniform of degree k, then it must have "local correlation" with a function of the form $\omega^{\phi(x)}$, where $\omega = \exp 2\pi i/N$ and ϕ is a polynomial of degree d. "Local" in this context means that one can partition $\mathbb{Z}_N$ into arithmetic progressions of size N^η (for some η that depends on c and k only) on a large proportion of which one can find such a correlation.

This was strong enough to prove Szemerédi's theorem, but for several other applications the highly local nature of the correlation is too weak. However, in the quadratic case, this problem has been remedied by Green and Tao [GT05a]. In this case, the obstacle to "globalizing" the argument is that a certain globally-defined bilinear form that occurs in the proof of [Gow01] is not symmetric, and thus does not allow one to define a corresponding globally-defined quadratic form. (In the context of $\mathbb{Z}_N$, "global" means something like "defined on a proportional-sized Bohr set". For $\mathbb{F}_p^n$ one can take it to mean "defined everywhere".) Green and Tao discovered an ingenious "symmetry argument" that allows one to replace the bilinear form by one that *is* symmetric, and this allowed them to prove a quadratic structure theorem for functions with large U^3-norm that is closely analogous to the linear structure theorem that follows from conventional Fourier analysis.

An excellent exposition of this structure theorem when the group G is a vector space over a finite field can be found in [Gre05b]. This contains proofs of all the background results that we state here.

Recall that in the linear case, we called f_1 "structured" because it was constant on affine subspaces of low codimension. For quadratic Fourier analysis, we need a quadratic analogue of the notion of a decomposition of $\mathbb{F}_p^n$ into parallel affine subspaces of codimension d_1. In order to define such a decomposition, one can take a surjective linear map $\Gamma_1 : \mathbb{F}_p^n \to \mathbb{F}_p^{d_1}$ and for each $a \in \mathbb{F}_p^{d_1}$ one can set V_a to equal $\Gamma_1^{-1}(\{a\})$. If we want to make this idea quadratic, we should replace the linear map Γ_1 by a "quadratic map" Γ_2, which we do in a natural way as follows. We say that a function $\Gamma_2 : \mathbb{F}_p^n \to \mathbb{F}_p^{d_2}$ is *quadratic* if it is of the form $x \mapsto (q_1(x), \dots, q_{d_2}(x))$, where $q_1, \dots, q_{d_2}$ are quadratic forms on $\mathbb{F}_p^n$. Then, for each $b \in \mathbb{F}_p^{d_2}$ we define W_b to be $\{x \in \mathbb{F}_p^n : \Gamma_2(x) = b\}$.

In [GT05b], Green and Tao define $\mathcal{B}_1$ to be the algebra generated by the sets V_a and $\mathcal{B}_2$ for the finer algebra generated by the sets $V_a \cap W_b$. They call $\mathcal{B}_1$ a *linear factor of complexity* d_1 and $(\mathcal{B}_1, \mathcal{B}_2)$ a *quadratic factor of complexity* (d_1, d_2). This is to draw out a close analogy with the "characteristic factors" that occur in ergodic theory.

These definitions give us a suitable notion of a "quadratically structured" function—it is a function f_1 for which we can find a linear map $\Gamma_1 : \mathbb{F}_p^n \to \mathbb{F}_p^{d_1}$ and a quadratic map $\Gamma_2 : \mathbb{F}_p^n \to \mathbb{F}_p^{d_2}$ such that d_1 and d_2 are not too large and f_1 is constant on the sets $V_a \cap W_b$ defined above. This is equivalent to saying that f_1 is measurable with respect to the algebra $\mathcal{B}_2$, and also to saying that $f_1(x)$ depends on $(\Gamma_1(x), \Gamma_2(x))$ only.

The quadratic structure theorem of Green and Tao implies that a bounded function f defined on $\mathbb{F}_p^n$ can be written as a sum $f_1 + f_2$, where f_1 is quadratically structured in the above sense, and $\|f_2\|_{U^3}$ is small. In [GT05b] the result is stated explicitly for $p = 5$, but this is merely because of the emphasis placed on 4-term progressions. The proof is not affected by the choice of p (as long as it stays fixed).

In the statement below, we write $\mathbb{E}(f|\mathcal{B}_2)$ for the conditional expectation, or averaging projection, of f. That is, if $X = V_a \cap W_b$ is an atom of $\mathcal{B}_2$ and $x \in X$, then $\mathbb{E}(f|\mathcal{B}_2)(x)$ is the average of f over X. Since the function $\mathbb{E}(f|\mathcal{B}_2)$ is constant on the sets $V_a \cap W_b$, it is quadratically structured in the sense that interests us.

Theorem 3.3. [GT05b] *Let p be a fixed prime, let $\delta > 0$ and suppose that $n > n_0(\delta)$ is sufficiently large. Given any function $f : \mathbb{F}_p^n \to [-1, 1]$, there exists a quadratic factor $(\mathcal{B}_1, \mathcal{B}_2)$ of complexity at most $((4\delta^{-1})^{3C_0+1}, (4\delta^{-1})^{2C_0+1})$ together with a decomposition*

$$f = f_1 + f_2,$$

where

$$f_1 := \mathbb{E}(f|\mathcal{B}_2) \quad \textit{and} \quad \|f_2\|_{U^3} \leq \delta.$$

The absolute constant C_0 can be taken to be 2^{16}.

As it stands, the above theorem is not quite suitable for applications, because technical problems arise if one has to deal with quadratic forms of low rank. (Notice that so far we have said nothing about the quadratic forms q_i—not even that they are distinct.) Let $\Gamma_2 = (q_1, \dots, q_k)$ be a quadratic map and for each i let β_i be the symmetric bilinear form corresponding to q_i: that is, $\beta_i(x, y) = (q_i(x+y) - q_i(x) - q_i(y))/2$. We shall say that Γ_2 *is of rank at least r* if the bilinear form $\sum_i \lambda_i \beta_i$ has rank at least r whenever $\lambda_1, \dots, \lambda_{d_2}$ are elements of $\mathbb{F}_p$ that are not all zero. If Γ_2 is used in combination with some linear map Γ_1 to define a quadratic factor $(\mathcal{B}_1, \mathcal{B}_2)$, then we shall also say that this quadratic factor has rank at least r.

Just to clarify this definition, let us prove a simple lemma that will be used later.

Lemma 3.4. *Let β be a symmetric bilinear form of rank r on $\mathbb{F}_p^n$ and let W be a subspace of $\mathbb{F}_p^n$ of codimension d_1. Then the rank of the restriction of β to W is at least $r - 2d_1$.*

Proof. Let $V = \mathbb{F}_p^n$. For every subspace W of V, let us write $W^\perp$ for the subspace

$$\{v \in V : \beta(v, w) = 0 \text{ for every } w \in W\}.$$

Let us define the *nullity* of β to be the dimension of $V^\perp$. Then the rank of β is equal to n minus its nullity, which is the codimension of $V^\perp$. We are assuming that this is r.

Now let W have codimension d_1. We begin by bounding from above the dimension of $W^\perp$. To do this, let Y be a complement for W, which, by hypothesis, will have dimension d_1. Then $V^\perp = W^\perp \cap Y^\perp$ and $Y^\perp$ has dimension at least $n - d_1$, so the dimension of $W^\perp$ is at most $d_1 + \dim(V^\perp)$, which, by hypothesis, is at least $d_1 + n - r$. Therefore, the codimension of $W^\perp$ is at most $r - d_1$, which implies that the codimension of $W^\perp$ inside W is at most $r - 2d_1$. This implies the result. □

We are now in a position to state the version of the structure theorem that we shall be using. It can be read out of (but is not explicitly stated in) [Gre05b] and [GT05b].

Theorem 3.5. *Let p be a fixed prime, let $\delta > 0$, let $r : \mathbb{N} \to \mathbb{N}$ an arbitrary function (which may depend on δ) and suppose that $n > n_0(r, \delta)$ is sufficiently large. Then given any function $f : \mathbb{F}_p^n \to [-1, 1]$, there exists $d_0 = d_0(r, \delta)$ and a quadratic factor $(\mathcal{B}_1, \mathcal{B}_2)$ of rank at least $r(d_1 + d_2)$ and complexity at most $(d_1, d_2), d_1, d_2 \leq d_0$, together with a decomposition*

$$f = f_1 + f_2 + f_3,$$

where

$$f_1 := \mathbb{E}(f|\mathcal{B}_2), \quad \|f_2\|_2 \leq \delta \quad \text{and} \quad \|f_3\|_{U^3} \leq \delta.$$

Note that $\mathbb{E}f_1 = \mathbb{E}f$. In particular $\mathbb{E}f_1 = 0$ whenever f is the balanced function of a subset of $\mathbb{F}_p^n$. It can be shown that f_1 is uniform whenever f is uniform: roughly speaking, the reason for this is that $\mathbb{E}(f|\mathcal{B}_1)$ is approximately zero and the atoms of $\mathcal{B}_2$ are uniform subsets of the atoms of $\mathcal{B}_1$. However, we shall not need this fact.

We shall apply Theorem 3.5 when r is the function $d \mapsto 2md + C$ for a constant C. Unfortunately, ensuring that factors have high rank is an expensive process: even for this modest function the argument involves an iteration that increases d_0 exponentially at every step. For this reason we have stated the theorem in a qualitative way. A quantitative version would involve a tower-type bound.

3.3. **Square-independence is sufficient.** We now have the tools we need to show that square-independence coupled with CS-complexity 2 is sufficient to guarantee the correct number of solutions in uniform sets. The basic idea of the proof is as follows. Given a set $A \subset \mathbb{F}_p^n$ of density α, we first replace it by its balanced function $f(x) = A(x) - \alpha$. Given a square-independent linear system $\mathcal{L}$ of complexity at most 2, our aim is to show, assuming that $\|f\|_{U^2}$ is sufficiently small, that

$$\mathbb{E}_{\mathbf{x}\in(\mathbb{F}_p^n)^d}\prod_{i=1}^m f(L_i(\mathbf{x}))$$

is also small. (Once we have done that, it will be straightforward to show that the same average, except with A replacing f, is close to α^m.) In order to carry out this estimate, we first apply the structure theorem to decompose f as $f_1 + f_2 + f_3$, where f_1 is quadratically structured, f_2 is small in L_2 and f_3 is quadratically uniform. This then allows us to decompose the product into a sum of 3^m products, one for each way of choosing f_1, f_2 or f_3 from each of the m brackets. If we ever choose f_2, then the Cauchy-Schwarz inequality implies that the corresponding term is small, and if we ever choose f_3 then a similar conclusion follows from Theorem 2.3. Thus, the most important part of the proof is to use the linear uniformity and quadratic structure of f_1 to prove that the product

$$\mathbb{E}_{\mathbf{x}\in(\mathbb{F}_p^n)^d}\prod_{i=1}^m f_1(L_i(\mathbf{x}))$$

is small. This involves a calculation that generalizes the one we used to prove Theorem 3.1. The main step is the following lemma, where we do the calculation in the case where the linear factor $\mathcal{B}_1$ is trivial.

Lemma 3.6. *Let $\mathcal{L} = (L_1, \dots, L_m)$ be a square-independent system of linear forms and let $\Gamma_2 = (q_1, \dots, q_{d_2})$ be a quadratic map from $\mathbb{F}_p^n$ to $\mathbb{F}_p^{d_2}$ of rank at least r. Let $\phi_1, \dots, \phi_m$ be linear maps from $(\mathbb{F}_p^n)^d$ to $\mathbb{F}_p^{d_2}$ and let $b_1, \dots, b_m$ be elements of $\mathbb{F}_p^{d_2}$. Let $\mathbf{x} = (x_1, \dots, x_d)$ be a randomly chosen element of $(\mathbb{F}_p^n)^d$. Then the probability that $\Gamma_2(L_i(\mathbf{x})) = \phi_i(\mathbf{x}) + b_i$ for every i differs from p^{-md_2} by at most $p^{-r/2}$.*

Proof. Let Λ be the set of all $m \times d_2$ matrices $\lambda = (\lambda_{ij})$ over $\mathbb{F}_p$ and let us write $\phi_i = (\phi_{i1}, \dots, \phi_{id_2})$ and $b_i = (b_{i1}, \dots, b_{id_2})$ for each i. The probability we are interested in is the probability that $q_j(L_i(\mathbf{x})) = \phi_{ij}(\mathbf{x}) + b_{ij}$ for every $i \le m$ and every $j \le d_2$. This equals

$$\mathbb{E}_{\mathbf{x}}\mathbb{E}_{\lambda\in\Lambda}\prod_{i=1}^m\prod_{j=1}^{d_2}\omega^{\lambda_{ij}(q_j(L_i(\mathbf{x}))-\phi_{ij}(\mathbf{x})-b_{ij})},$$

since if $q_j(L_i(\mathbf{x})) = \phi_{ij}(\mathbf{x}) + b_{ij}$ for every i and j, then the expectation over λ is 1, and otherwise if we choose i and j such that $q_j(L_i(\mathbf{x})) \neq \phi_{ij}(\mathbf{x})+b_{ij}$ and consider the expectation over λ_{ij} while all other entries of λ are fixed, then we see that the expectation over λ is zero.

We can rewrite the above expectation as

$$\mathbb{E}_{\lambda\in\Lambda}\mathbb{E}_{\mathbf{x}}\omega^{\sum_{i,j}\lambda_{ij}(q_j(L_i(\mathbf{x}))-\phi_{ij}(\mathbf{x})-b_{ij})}.$$

If $\lambda = 0$, then obviously the expectation over $\mathbf{x}$ is 1. This happens with probability p^{-md_2}. Otherwise, for each i let us say that the coefficients of L_i are $c_{i1},\dots,c_{id}$. That is, let $L_i(\mathbf{x}) = \sum_{u=1}^d c_{iu}x_u$. Then

$$q_j(L_i(\mathbf{x})) = \sum_{u,v} c_{iu}c_{iv}\beta_j(x_u,x_v),$$

where β_j is the bilinear form associated with q_j. Choose some j such that λ_{ij} is non-zero for at least one i. Then the square-independence of the linear forms L_i implies that there exist u and v such that $\sum_i \lambda_{ij}c_{iu}c_{iv}$ is not zero.

Fix such a j, u and v and do it in such a way that $u = v$, if this is possible. We shall now consider the expectation as x_u and x_v vary with every other x_w fixed. Notice first that

$$\sum_{i,j}\lambda_{ij}q_j(L_i(\mathbf{x})) = \sum_{i,j}\sum_{t,w}\lambda_{ij}c_{it}c_{iw}\beta_j(x_t,x_w).$$

Let us write β_{tw} for the bilinear form $\sum_{i,j}\lambda_{ij}c_{it}c_{iw}\beta_j$, so that this becomes $\sum_{t,w}\beta_{tw}(x_t,x_w)$. Let us also write $\phi(\mathbf{x})$ for $\sum_{ij}\lambda_{ij}\phi_{ij}(\mathbf{x})$ and let $\phi_1,\dots,\phi_d$ be linear maps from $\mathbb{F}_p^n$ to $\mathbb{F}_p$ such that $\phi(\mathbf{x}) = \sum_t \phi_t(x_t)$ for every $\mathbf{x}$. Then

$$\sum_{i,j}\lambda_{ij}(q_j(L_i(\mathbf{x})) - \phi_{ij}(\mathbf{x})) = \sum_{t,w}\beta_{tw}(x_t,x_w) - \sum_t \phi_t(x_t).$$

Notice that if we cannot get u to equal v, then $\sum_i \lambda_{ij}c_{iu}^2 = 0$ for every u and every j, which implies that $\beta_{uu} = 0$. Notice also that the assumption that Γ_2 has rank at least r and the fact that $\sum_i \lambda_{ij}c_{iu}c_{iw} \neq 0$ for at least one j imply that β_{uv} has rank at least r.

If we fix every x_t except for x_u and x_v, then $\sum_{t,w}\beta_{tw}(x_t,x_w) - \sum_t \phi_t(x_t)$ is a function of x_u and x_v of the form

$$\beta_{uv}(x_u,x_v) + \psi_u(x_u) + \psi_v(x_v),$$

where ψ_u and ψ_v are linear functionals on $\mathbb{F}_p^n$ (that depend on the other x_t).

Now let us estimate the expectation

$$\mathbb{E}_{x_u,x_v}\omega^{\sum_{i,j}\lambda_{ij}(q_j(L_i(\mathbf{x}))-\phi_{ij}(\mathbf{x})-b_{ij})},$$

where we have fixed every x_t apart from x_u and x_v. Letting $b = \sum \lambda_{ij} b_{ij}$ and using the calculations we have just made, we can write this in the form

$$\mathbb{E}_{x_u, x_v} \omega^{\beta_{uv}(x_u, x_v) + \psi_u(x_u) + \psi_v(x_v) - b}.$$

If $u = v$, then the expectation is just over x_u and the exponent has the form $q(x_u) + w^T u - b$ for some quadratic form q of rank at least r. Therefore, by Lemma 3.2, the expectation is at most $p^{-r/2}$. If $u \neq v$ (and therefore every $\mathbf{b}_{uu}$ is zero) then for each x_v the exponent is linear in u. This means that either the expectation over x_u is zero or the function $\beta_{uv}(x_u, x_v) + \psi_u(x_u)$ is constant. If the latter is true when $x_v = y$ and when $x_v = z$, then $\beta_{uv}(x_u, y - z)$ is also constant, and therefore identically zero. Since β_{uv} has rank at least r, $y - z$ must lie in a subspace of codimension at least r. Therefore, the set of x_v such that $\beta_{uv}(x_u, x_v) + \psi_u(x_u)$ is constant is an affine subspace of $\mathbb{F}_p^n$ of codimension at least r, which implies that the probability (for a randomly chosen x_v) that the expectation (over x_u) is non-zero is at most p^{-r}. When the expectation is non-zero, it has modulus 1.

In either case, we find that, for any non-zero $\lambda \in \Lambda$, the expectation over $\mathbf{x}$ is at most $p^{-r/2}$, and this completes the proof of the lemma. □

We now want to take into account Γ_1 as well as Γ_2. This turns out to be a short deduction from the previous result. First let us do a simple piece of linear algebra.

Lemma 3.7. *Let $\mathcal{L} = (L_1, \dots, L_m)$ be a collection of linear forms in d variables, and suppose that the linear span of $L_1, \dots, L_m$ has dimension d'. Let $\Gamma_1 : \mathbb{F}_p^n \to \mathbb{F}_p^{d_1}$ be a surjective linear map and let $\phi : (\mathbb{F}_p^n)^d \to (\mathbb{F}_p^{d_1})^m$ be defined by the formula*

$$\phi : \mathbf{x} \mapsto (\Gamma_1(L_1(\mathbf{x})), \dots, \Gamma_1(L_m(\mathbf{x}))).$$

Then the image of ϕ is the subspace Z of $(\mathbb{F}_p^{d_1})^m$ that consists of all sequences $(a_1, \dots, a_m)$ such that $\sum_i \mu_i a_i = 0$ whenever $\sum_i \mu_i L_i = 0$. The dimension of Z is $d'd_1$.

Proof. Since the m forms L_i span a space of dimension d', the set of sequences $\mu = (\mu_1, \dots, \mu_m)$ such that $\sum_i \mu_i L_i = 0$ is a linear subspace W of $\mathbb{F}_p^m$ of dimension $m - d'$. Therefore, the condition that $\sum_i \mu_i a_i = 0$ for every sequence $\mu \in W$ restricts $(a_1, \dots, a_m)$ to a subspace of $(\mathbb{F}_p^{d_1})^m$ of codimension $d_1(m - d')$. (An easy way to see this is to write $a_i = (a_{i1}, \dots, a_{id_1})$ and note that for each j the sequence $(a_{1j}, \dots, a_{mj})$ is restricted to a subspace of codimension $m - d'$.) Therefore, the dimension of Z is $d'd_1$, as claimed.

Now let us show that Z is the image of ϕ. Since ϕ is linear, Z certainly contains the image of ϕ, so it will be enough to prove that the rank of ϕ is $d'd_1$.

Abusing notation, let us write $\Gamma_1(\mathbf{x})$ for the sequence $(\Gamma_1 x_1, \dots, \Gamma_1 x_d)$, which belongs to $(\mathbb{F}_p^{d_1})^d$. Then $\phi(\mathbf{x})$ can be rewritten as $(L_1(\Gamma_1(\mathbf{x})), \dots, L_m(\Gamma_1(\mathbf{x})))$. Since Γ_1 is a surjection, it is also a surjection when considered as a map on $(\mathbb{F}_p^n)^d$. Therefore, the rank of ϕ is the rank of the map $\psi : (\mathbb{F}_p^{d_1})^d \to (\mathbb{F}_p^{d_1})^m$ defined by

$$\psi : \mathbf{y} \mapsto (L_1(\mathbf{y}), \dots, L_m(\mathbf{y})).$$

Since the L_i span a space of dimension d', the nullity of this map is $d_1(d - d')$, so its rank is $d_1 d'$. Therefore, the image of ϕ is indeed Z. □

Lemma 3.8. *Let $\mathcal{L} = (L_1, \dots, L_m)$ be a square-independent system of linear forms in d variables, and suppose that the linear span of $L_1, \dots, L_m$ has dimension d'. Let $\Gamma_1 : \mathbb{F}_p^n \to \mathbb{F}_p^{d_1}$ be a surjective linear map and let $\Gamma_2 : \mathbb{F}_p^n \to \mathbb{F}_p^{d_2}$ be a quadratic map of rank at least r. Let $a_1, \dots, a_m$ be elements of $\mathbb{F}_p^{d_1}$ and let $b_1, \dots, b_m$ be elements of $\mathbb{F}_p^{d_2}$, and let ϕ and Z be as defined in the previous lemma. Then the probability, if $\mathbf{x}$ is chosen randomly from $(\mathbb{F}_p^n)^d$, that $\Gamma_1(L_i(\mathbf{x})) = a_i$ and $\Gamma_2(L_i(\mathbf{x})) = b_i$ for every $i \le m$ is zero if $(a_1, \dots, a_m) \in Z$, and otherwise it differs from $p^{-d_1 d' - d_2 m}$ by at most $p^{d_1 - d' d_1 - r/2}$.*

Proof. If $\mathbf{a} = (a_1, \dots, a_m) \notin Z$, then there exists $\mu \in \mathbb{F}_p^m$ such that $\sum_i \mu_i a_i \neq 0$ and $\sum_i \mu_i L_i(\mathbf{x}) = 0$ for every $\mathbf{x}$. Since Γ_1 is linear, it follows that there is no $\mathbf{x}$ such that $\Gamma_1(L_i(\mathbf{x})) = a_i$ for every i.

Otherwise, by Lemma 3.7, $\mathbf{a}$ lies in the image of ϕ, which has rank $d' d_1$, so $\phi^{-1}(\{\mathbf{a}\})$ is an affine subspace of $(\mathbb{F}_p^n)^d$ of codimension $d' d_1$. Therefore, the probability that $\phi(\mathbf{x}) = \mathbf{a}$ is $p^{-d' d_1}$. Now let us use Lemma 3.6 to estimate the probability, conditional on this, that $\Gamma_2(L_i(\mathbf{x})) = b_i$ for every $i \le m$.

In the proof of Lemma 3.7, we observed that $\phi(\mathbf{x})$ depends on $\Gamma_1(\mathbf{x})$ only, so we shall estimate the required probability, given the value of $\Gamma_1(\mathbf{x})$. (Recall that this is notation for $(\Gamma_1 x_1, \dots, \Gamma_1 x_d)$.) In order to specify the set on which we are conditioning, let V be the kernel of Γ_1 (considered as a map defined on $\mathbb{F}_p^n$), and given a sequence $(w_1, \dots, w_d) \in (\mathbb{F}_p^n)^d$, let us estimate the required probability, given that $x_u \in V + w_u$ for every u.

Let us write $x_u = y_u + w_u$. Thus, we are estimating the probability that $\Gamma_2(L_i(\mathbf{y} + \mathbf{w})) = b_i$ for every $i \le m$. But for each i we can write $\Gamma_2(L_i(\mathbf{y} + \mathbf{w}))$ as $\Gamma_2(L_i(\mathbf{y})) + \phi_i(\mathbf{y}) + b_i'$ for some linear function $\phi_i : V^d \to \mathbb{F}_p^{d_2}$ and some vector $b_i' \in \mathbb{F}_p^{d_2}$.

Because Γ_2 has rank at least r and the codimension of V in $\mathbb{F}_p^n$ is d_1, Lemma 3.4 implies that the rank of the restriction of Γ_2 to V is at least $r - 2d_1$. Therefore, by Lemma 3.6, the probability that $\Gamma_2(L_i(\mathbf{y})) = -\phi_i(\mathbf{y}) + b_i - b_i'$ for every i differs from $p^{-m d_2}$ by at most $p^{d_1 - r/2}$.

Since this is true for all choices of $\mathbf{w}$, we have the same estimate if we condition on the event that $\phi(\mathbf{x}) = \mathbf{a}$ for some fixed $\mathbf{a} \in Z$. Therefore, the probability that $\Gamma_1(L_i(\mathbf{x})) = a_i$ and $\Gamma_2(L_i(\mathbf{x})) = b_i$ for every i differs from $p^{-d'd_1 - md_2}$ by at most $p^{d_1 - d'd_1 - r/2}$, as claimed. □

Next, we observe that Lemma 3.8 implies that all the atoms of $\mathcal{B}_2$ have approximately the same size.

Corollary 3.9. *Let Γ_1 and Γ_2 be as above and let x be a randomly chosen element of $\mathbb{F}_p^n$. Then for every $a \in \mathbb{F}_p^{d_1}$ and every $b \in \mathbb{F}_p^{d_2}$, the probability that $\Gamma_1(x) = a$ and $\Gamma_2(x) = b$ differs from $p^{-d_1 - d_2}$ by at most $p^{-r/2}$.*

Proof. Let us apply Lemma 3.8 in the case where $\mathcal{L}$ consists of the single one-variable linear form $L(x) = x$. This has linear rank 1 and is square-independent, so when we apply the lemma we have $d' = m = 1$. If we let $a_1 = a$ and $b_1 = b$, then the conclusion of the lemma tells us precisely what is claimed. □

The next two lemmas are simple technical facts about projections on to linear factors. The first one tells us that if g is any function that is uniform and constant on the atoms of a linear factor, then it has small L_2-norm. The second tells us that projecting on to a linear factor decreases the U^2-norm.

Lemma 3.10. *Let G be a function from $\mathbb{F}_p^{d_1}$ to $[-1, 1]$, let $\Gamma_1 : \mathbb{F}_p^n \to \mathbb{F}_p^{d_1}$ be a surjective linear map and let $g = G \circ \Gamma_1$. Then $\|g\|_2^4 \leq p^{d_1} \|g\|_{U^2}^4$.*

Proof. Since Γ_1 takes each value in $\mathbb{F}_p^{d_1}$ the same number of times, $\|g\|_{U^2} = \|G\|_{U^2}$. But

$$\|G\|_{U^2}^4 = \mathbb{E}_a(\mathbb{E}_b G(b)G(b+a))^2 \geq p^{-d_1}(\mathbb{E}_b G(b)^2)^2 = p^{-d_1}\|G\|_2^4,$$

which proves the result, since $\|g\|_2 = \|G\|_2$ as well. □

Lemma 3.11. *Let f be a function from $\mathbb{F}_p^n$ to $\mathbb{R}$, let $\mathcal{B}_1$ be a linear factor and let $g = \mathbb{E}(f|\mathcal{B}_1)$. Then $\|g\|_{U^2} \leq \|f\|_{U^2}$.*

Proof. On every atom of $\mathcal{B}_1$, g is constant and $f - g$ averages zero. Let Γ_1 be the linear map that defines $\mathcal{B}_1$ and, as we did earlier, for each $a \in \mathbb{F}_p^{d_1}$ let V_a stand for $\Gamma_1^{-1}(\{a\})$. Then

$$\|f\|_{U^2}^4 = \mathbb{E}_{a_1+a_2=a_3+a_4}\mathbb{E}_{\substack{x_1+x_2=x_3+x_4 \\ \Gamma_1(x_i)=a_i}} f(x_1)f(x_2)f(x_3)f(x_4)\ .$$

Let us fix a choice of $a_1+a_2 = a_3+a_4$ and consider the inner expectation. Setting $g' = f-g$, this has the form

$$\mathbb{E}_{\substack{x_1+x_2=x_3+x_4 \\ \Gamma_1(x_i)=a_i}} (\lambda_1 + g'(x_1))(\lambda_2 + g'(x_2))(\lambda_3 + g'(x_3))(\lambda_4 + g'(x_4))$$

This splits into sixteen parts. Each part that involves at least one λ_i and at least one $g'(x_i)$ is zero, because any three of the x_is can vary independently and g' averages zero on every atom of $\mathcal{B}_1$. This means that the expectation is

$$\lambda_1\lambda_2\lambda_3\lambda_4 + \mathbb{E}_{\substack{x_1+x_2=x_3+x_4\\ \Gamma_1(x_i)=a_i}} g'(x_1)g'(x_2)g'(x_3)g'(x_4)\ .$$

If we now take expectations over $a_1+a_2=a_3+a_4$ we find that $\|f\|_{U^2}^4 = \|g\|_{U^2}^4 + \|f-g\|_{U^2}^4$. Notice that this is a general result about how the U^2-norm of a function is related to the U^2-norm of a projection on to a linear factor. □

Now we are ready to estimate the product we are interested in, for functions that are constant on the atoms of $\mathcal{B}_2$.

Lemma 3.12. *Let $\Gamma_1 : \mathbb{F}_p^n \to \mathbb{F}_p^{d_1}$ be a linear function and $\Gamma_2 : \mathbb{F}_p^n \to \mathbb{F}_p^{d_2}$ be a quadratic function. Let $(\mathcal{B}_1, \mathcal{B}_2)$ be the corresponding quadratic factor and suppose that this has rank at least r. Let $c > 0$ and let $f : \mathbb{F}_p^n \to [-1,1]$ be a function with $\|f\|_{U^2} \le c$ and let $f_1 = \mathbb{E}(f|\mathcal{B}_2)$. Let $\mathcal{L} = (L_1, \dots, L_m)$ be a square-independent system of linear forms. Then*

$$\mathbb{E}_{\mathbf{x}\in(\mathbb{F}_p^n)^d} \prod_{i=1}^m f_1(L_i(\mathbf{x})) \le 4^m c p^{d_1/4} + 2^{m+1} p^{m(d_1+d_2)-r/2}.$$

Proof. Let $g = \mathbb{E}(f_1|\mathcal{B}_1)$ and let $h = f_1 - g$. Then $\|g\|_1 \le \|g\|_2 \le p^{d_1/4}\|g\|_{U^2}$, by Cauchy-Schwarz and Lemma 3.10. By Lemma 3.11, $\|g\|_{U^2} \le \|f\|_{U^2}$, which is at most c, by hypothesis. Therefore, $\|g\|_1 \le cp^{d_1/4}$.

Since $f_1 = g+h$, we can split the product up into a sum of 2^m products, in each of which we replace $f_1(L_i(\mathbf{x}))$ by either $g(L_i(\mathbf{x}))$ or $h(L_i(\mathbf{x}))$. Since $\|g\|_1 \le cp^{d_1/4}$ and $\|h\|_\infty \le 2$, any product that involves at least one g has average at most $2^m cp^{d_1/4}$. It remains to estimate

$$\mathbb{E}_{\mathbf{x}\in(\mathbb{F}_p^n)^d} \prod_{i=1}^m h(L_i(\mathbf{x})).$$

Let Z be as defined in Lemma 3.7, and for each $\mathbf{a} = (a_1, \dots, a_m)$ and $\mathbf{b} = (b_1, \dots, b_m)$, let $P(\mathbf{a}, \mathbf{b})$ be the probability that $\Gamma_1(L_i(\mathbf{x})) = a_i$ and $\Gamma_2(L_i(\mathbf{x})) = b_i$ for every i. By Lemma 3.8, we can set $P(\mathbf{a}, \mathbf{b}) = p^{-d'd_1 - md_2} + \epsilon(\mathbf{a}, \mathbf{b})$, with $|\epsilon(\mathbf{a}, \mathbf{b})| \le p^{d_1 - d'd_1 - r/2}$.

Now let H be defined by the formula $h(x) = H(\Gamma_1 x, \Gamma_2 x)$. Because h is constant on the atoms of $\mathcal{B}_2$, H is well-defined on the set of all elements of $\mathbb{F}_p^{d_1} \times \mathbb{F}_p^{d_2}$ of the form $(\Gamma_1 x, \Gamma_2 x)$. Since h takes values in $[-2, 2]$, so does H.

Next, we show that $\mathbb{E}_b H(a, b)$ is small for any fixed $a \in \mathbb{F}_p^{d_1}$, using the facts that h averages 0 on every cell of $\mathcal{B}_1$ and that it is constant on the cells of $\mathcal{B}_2$. Let us fix an a

and write $P(b)$ for the probability that $\Gamma_2(x) = b$ given that $\Gamma_1(x) = a$—that is, for the density of $V_a \cap W_b$ inside V_a. Then

$$0 = \mathbb{E}_{x \in V_a} h(x) = \mathbb{E}_{x \in V_a} H(\Gamma_1 x, \Gamma_2 x) = \sum_b P(b) H(a, b).$$

By Corollary 3.9, we can write $P(b) = p^{-d_2} + \epsilon(b)$, with $|\epsilon(b)| \leq p^{d_1 - r/2}$ for every b. Therefore, the right-hand side differs from $\mathbb{E}_b H(a, b)$ by at most $2p^{d_1+d_2-r/2}$, which implies that $|\mathbb{E}_b H(a, b)| \leq 2p^{d_1+d_2-r/2}$.

Now

$$\mathbb{E}_{\mathbf{x}} \prod_{i=1}^m h(L_i(\mathbf{x})) = \mathbb{E}_{\mathbf{x}} \prod_{i=1}^m H(\Gamma_1(L_i(\mathbf{x})), \Gamma_2(L_i(\mathbf{x}))) = \sum_{\mathbf{a} \in Z} \sum_{\mathbf{b}} P(\mathbf{a}, \mathbf{b}) \prod_{i=1}^m H(a_i, b_i).$$

Let us split up this sum as

$$p^{-d' d_1 - m d_2} \sum_{\mathbf{a} \in Z} \sum_{\mathbf{b}} \prod_{i=1}^m H(a_i, b_i) + \sum_{\mathbf{a} \in Z} \sum_{\mathbf{b}} \epsilon(\mathbf{a}, \mathbf{b}) \prod_{i=1}^m H(a_i, b_i).$$

The first term equals $\mathbb{E}_{a \in Z} \prod_{i=1}^m (\mathbb{E}_b H(a_i, b))$, which is at most $(2p^{d_1+d_2-r/2})^m$. The second is at most $p^{(d' d_1 + m d_2)} 2^m p^{d_1 - d' d_1 - r/2} = 2^m p^{d_1 + m d_2 - r/2}$. Therefore, the whole sum is at most $2^{m+1} p^{m(d_1+d_2)-r/2}$. Together with our estimate for the terms that involved g, this proves the lemma. □

We have almost finished the proof of our main result.

Theorem 3.13. *For every $\epsilon > 0$ there exists $c > 0$ with the following property. Let $f : \mathbb{F}_p^n \to [-1, 1]$ be a c-uniform function. Let $\mathcal{L} = (L_1, \dots, L_m)$ be a square-independent system of linear forms in d variables, with Cauchy-Schwarz complexity at most 2. Then*

$$\mathbb{E}_{\mathbf{x} \in (\mathbb{F}_p^n)^d} \prod_{i=1}^m f(L_i(\mathbf{x})) \leq \epsilon.$$

Proof. Let $\delta > 0$ be a constant to be chosen later. Let C be such that $2^{m+1} p^{-C/2} \leq \epsilon/3$ and let r be the function $d \mapsto 2md + C$. Then according to the structure theorem (Theorem 3.5) there exists d_0, depending on r and δ only, and a quadratic factor $(\mathcal{B}_1, \mathcal{B}_2)$ of rank at least $2m(d_1 + d_2) + C$ and complexity (d_1, d_2), with d_1 and d_2 both at most d_0, such that we can write $f = f_1 + f_2 + f_3$, with $f_1 = \mathbb{E}(f|\mathcal{B}_2)$, $\|f_2\|_2 \leq \delta$ and $\|f_3\|_{U^3} \leq \delta$.

Let us show that the sum does not change much if we replace $f(L_m(\mathbf{x}))$ by $f_1(L_m(\mathbf{x}))$. The difference is what we get if we replace $f(L_m(\mathbf{x}))$ by $f_2(L_m(\mathbf{x})) + f_3(L_m(\mathbf{x}))$. Now $\|f_2\|_1 \leq \|f_2\|_2$ and $\|f\|_\infty \leq 1$, so the contribution from the f_2 part is at most δ. As for the

f_3 part, since $\|f_3\|_{U^3} \leq \delta$ and $\|f\|_\infty \leq 1$, Theorem 2.3 tells us that the contribution is at most δ. Therefore, the total difference is at most $\delta + \delta \leq 2\delta$.

Now let us replace f by f_1 in the penultimate bracket. The same argument works, since $\|f_1\|_\infty \leq 1$. Indeed, we can carry on with this process, replacing every single f by f_1, and the difference we make will be at most $2m\delta$.

We are left needing to show that the product with every f replaced by f_1 is small. This is what Lemma 3.12 tells us. It gives us an upper bound of $4^m cp^{d_1/4} + 2^{m+1}p^{m(d_1+d_2)-r/2}$, where for r we can take $2m(d_1 + d_2) + C$. Therefore, the upper bound is $4^m cp^{d_0/4} + 2^{m+1}p^{-C/2}$, which, by our choice of C, is at most $4^m cp^{d_0/4} + \epsilon/3$.

To finish, let $\delta = \epsilon/6m$. This determines the value of d_0 and we can then set c to be $4^{-m}p^{-d_0/4}\epsilon/3$, which will be a function of ϵ only. □

Because of our use of Theorem 3.5, the bounds in the above result and in the corollary that we are about to draw from it are both very weak. However, we have been explicit about all the bounds apart from d_0, partly in order to make it clear how the parameters depend on each other and partly to demonstrate that our weak bound derives just from the weakness of d_0 in the structure theorem.

Corollary 3.14. *For every $\epsilon > 0$ there exists $c > 0$ with the following property. Let A be a c-uniform subset of $\mathbb{F}_p^n$ of density α. Let $\mathcal{L} = (L_1, \dots, L_m)$ be a square-independent system of linear forms in d variables, with Cauchy-Schwarz complexity at most 2. Let $\mathbf{x} = (x_1, \dots, x_d)$ be a random element of $(\mathbb{F}_p^n)^d$. Then the probability that $L_i(\mathbf{x}) \in A$ for every i differs from α^m by at most ϵ.*

Proof. We shall choose as our c the c that is given by the previous theorem when ϵ is replaced by $\epsilon/2^m$. Our assumption is then that we can write $A = \alpha + f$ for a c-uniform function f. The probability we are interested in is

$$\mathbb{E}_{\mathbf{x}\in(\mathbb{F}_p^n)^d} \prod_{i=1}^m A(L_i(\mathbf{x})),$$

which we can split into 2^m parts, obtained by replacing each occurrence of A either by α or by f.

For each part that involves at least one occurrence of f, we have a power of α multiplied by a product over some subsystem of $\mathcal{L}$. This subsystem will also be square-independent and have CS-complexity at most 2. Moreover, the number of linear forms will have decreased. Therefore, the previous theorem and our choice of c tell us that the contribution it makes is at most $\epsilon/2^m$. Therefore, the contribution from all such parts is at most ϵ.

The only remaining part is the one where every $A(L_i(\mathbf{x}))$ has been replaced by α, and that gives us the main term α^m. □

4. Concluding remarks

First, we remark that Corollary 3.14 allows us to deduce rather straightforwardly a Szemerédi-type theorem for square-independent systems of CS-complexity 2 which have the additional property that they are *translation-invariant*. That is, one can show that any sufficiently dense subset of $\mathbb{F}_p^n$ contains a configuration of the given type.

Without the result of the preceding section, establishing that any sufficiently dense subset contains a solution to systems of this type would require a quadratic argument of the form used by Green and Tao to prove Szemerédi's Theorem for progressions of length 4 in finite fields [GT05b]. This would involve obtaining density increases on quadratic subvarieties of $\mathbb{F}_p^n$, which then need to be linearized in a carefully controlled manner. Although it is certainly possible to adapt their argument in this way, for purely qualitative purposes it is much simpler to use the result that configurations of this type are governed by the U^2-norm, which allows one to produce a density increase on an affine subspace. The resulting argument is almost identical to the well-known argument for 3-term progressions [Mes95]. Translation invariance is needed because the subspace on which we find a density increment may be an affine and not a strictly linear one. (It is not hard to show that the result is false if the system is not translation invariant.) Recall that two examples of square-independent translation invariant systems of complexity 2 are the systems $(x, x+y, x+z, x+y+z, x+y-z, x+z-y)$ and $(x, x+a, x+b, x+c, x+a+b, x+a+c, x+b+c)$.

The second of these examples shows that our main result implies the following useful "Pythagorean theorem," which generalizes the much more straightforward fact that if a is a constant and f averages zero, then $\|a+f\|_{U^2}^4 = a^4 + \|f\|_{U^2}^4$.

Theorem 4.1. *For every $\epsilon > 0$ there exists $c > 0$ such that if f is a c-uniform function from $\mathbb{F}_p^n$ to $[-1,1]$, $a \in [-1,1]$ is a constant, and $g(x) = a + f(x)$ for every $x \in \mathbb{F}_p^n$, then $\big|\|g\|_{U^3}^8 - (a^8 + \|f\|_{U^3}^8)\big| \leq \epsilon$.*

We briefly sketch the proof: expanding out the definition of $\|a+f\|_{U^3}^8$ one obtains a sum of 2^8 terms, one of which gives a^8 (if you choose a from every bracket) and one of which gives $\|f\|_{U^3}^8$ (if you choose f from every bracket). All the remaining terms are constant multiples of expectations of f over linear configurations that are square-independent and therefore, by our main result, small.

There are several ways in which the results of Section 3 might be generalized. An obvious one is to prove comparable results for the group $\mathbb{Z}_N$. As we mentioned earlier, we have a different proof for $\mathbb{F}_p^n$ and this can be transferred to $\mathbb{Z}_N$ by "semi-standard" methods. (That is, the general approach is clear, but the details can be complicated and sometimes require more than merely technical thought.) The alternative proof for $\mathbb{F}_p^n$ gives a doubly exponential bound for the main result rather than the tower-type bound obtained here.

Possibly even more obvious is to try to extend the main result of this paper to a proof of Conjecture 2.5. This involves a generalization in two directions: to systems of CS-complexity greater than 2, and to systems with true complexity greater than 1. All further cases will require polynomial Fourier analysis for a degree that is greater than 2: the simplest is likely to be to show that a square-independent system with CS-complexity 3 has true complexity 1. In this case, we would use a decomposition into a structured part (a projection onto a cubic factor) and a uniform part (which would be small in U^4 and therefore negligible) and then, as before, concentrate on the structured part. Square-independence (which implies cube-indepence) would ensure that we could reduce to the linear part of the factor as before.

This state of affairs leaves us very confident that Conjecture 2.5 is true. Although cubic and higher-degree Fourier analysis have not yet been worked out, they do at least exist in local form for $\mathbb{Z}_N$: they were developed in [Gow01] to prove the general case of Szemerédi's theorem. It is therefore almost certain that global forms will eventually become available, both for $\mathbb{Z}_N$ and for $\mathbb{F}_p^n$. And then, given a statement analogous to Theorem 3.5, it is easy to see how to generalize the main steps of our proof. In particular, the Gauss-sum estimates on which we depend so heavily have higher-degree generalizations.

A completely different direction in which one might consider generalizing the above results is to hypergraphs. For example, very similar proofs to those of Theorems 2.2 and 2.3 can be used to prove so-called "counting lemmas" for quasirandom hypergraphs—lemmas that assume that a certain norm is small and deduce that the hypergraph contains approximately the expected number of small configurations of a given kind.

One can now ask whether, as with sets, weaker quasirandomness assumptions about a hypergraph suffice to guarantee the right number of certain configurations, and if so, which ones. It turns out to be possible to give a complete answer to a fairly natural formulation of this question. Unfortunately, however, the proof is rather too easy to be interesting, so here we content ourselves with somewhat informal statements of results concerning the

special case of 3-uniform hypergraphs. The proofs we leave as exercises for any reader who might be interested.

Recall that if X, Y and Z are finite sets and $f : X \times Y \times Z \to \mathbb{R}$, then the *octahedral norm* of f is the eighth root of

$$\mathbb{E}_{x(0),x(1)\in X}\mathbb{E}_{y(0),y(1)\in Y}\mathbb{E}_{z(0),z(1)\in Z}\prod_{\epsilon\in\{0,1\}^3} f(x(\epsilon_1),y(\epsilon_2),z(\epsilon_3)).$$

It is easy to verify that if $X = Y = Z = G$ for some Abelian group G and $f(x,y,z) = g(x+y+z)$ for some function g, then the octahedral norm of f is the same as the U^3-norm of g. Therefore, it is natural to consider the octahedral norm of functions defined on $X \times Y \times Z$ as the correct analogue of the U^3-norm of functions defined on Abelian groups.

An important fact about the octahedral norm is that f has small octahedral norm if and only if it has a small correlation with any function of the form $u(x,y)v(y,z)w(x,z)$. Another important fact, the so-called "counting lemma" for quasirandom hypergraphs, states the following. Let X be a finite set and let H be a 3-uniform hypergraph with vertex set X and density α. Suppose that H is quasirandom in the sense that the function $H(x,y,z) - \alpha$ has small octahedral norm (where $H(x,y,z) = 1$ if $\{x,y,z\} \in H$ and 0 otherwise). Then H has about the expected number of copies of any fixed small hypergraph. For instance, if you choose x, y, z and w randomly from X, then the probability that all of $\{x,y,z\}, \{x,y,w\}, \{x,z,w\}$ and $\{y,z,w\}$ belong to H is approximately α^4.

Now let us suppose that g is uniform but not necessarily quadratically uniform, and that we again define $f(x,y,z)$ to be $g(x+y+z)$. What can we say about f? It is not necessarily the case that f has small octahedral norm, or that it has low correlation with functions of the form $u(x,y)v(y,z)w(x,z)$. However, it is not hard to show that it has low correlation with any function of the form $a(x)b(y)c(z)$, a property that was referred to as *vertex uniformity* in [Gow04].

One might therefore ask whether vertex uniformity was sufficient to guarantee the right number of copies of some small hypergraphs. However, well-known and easy examples shows that it does so only for hypergraphs such that no pair $\{x,y\}$ is contained in more than one hyperedge. For instance, let u be a random symmetric function from X^2 to $\{-1,1\}$ and let $H(x,y,z) = (3 + u(x,y) + u(y,z) + u(x,z))/6$. Then H is vertex uniform and has density $1/2$, but it is a simple exercise to show that $\mathbb{E}_{x,y,z,w}H(x,y,z)H(x,y,w)$ is about $5/18$ instead of the expected $1/4$.

However, this is perhaps not the right question to be asking. If g is uniform, then f has a stronger property than just vertex uniformity: one can prove that it does not correlate

with any function of the form $u(x,y)w(x,z)$, $u(x,y)v(y,z)$ or $v(y,z)w(x,z)$. If we take *this* as our definition of "weak quasirandomness" for functions (and call the hypergraph H weakly quasirandom if the function $H-\alpha$ is), then which hypergraphs appear with the right frequency (or with "frequency zero" if we are talking about functions rather than sets)? The answer turns out to be that a sum over copies of a small hypergraph H' will have the "right" value if and only if there is a pair $\{x,y\}$ that belongs to exactly one hyperedge $\{x,y,z\}$ of H'. The proof in the "if" direction is an easy exercise. In particular, it does not involve any interesting results about decomposing hypergraphs, which suggests that the main result of this paper is, in a certain sense, truly arithmetical.

As for the "only if" direction, here is a quick indication of how to produce an example (in the complex case, for simplicity). Suppose that no pair $\{x,y\}$ belongs to more than m hyperedges in H'. For each k between 2 and m let $f_k : X^2 \to \mathbb{C}$ be a function whose values are randomly chosen kth roots of unity. Then let $f(x,y,z)$ be the sum of all functions of the form $u(x,y)v(y,z)w(x,z)$, where each of u, v and w is some f_k with $2 \leq k \leq m$. When one expands out the relevant sum for this function f, one finds that most terms cancel, but there will be some that don't and they will all make a positive contribution. To find such a term, the rough idea is to choose for each face F of H' a triple of functions $(f_{k_1}, f_{k_2}, f_{k_3})$, where k_1, k_2 and k_3 are the number of faces of H' that include each of the three edges that make up the face F. For this term, each time a kth root of unity appears in the product, it is raised to the power k, so the term is large.

Finally, let us say just a little bit more about the result of Leibman mentioned at the beginning of the paper. The question in ergodic theory which is analogous to the one we were studying in Section 3 concerns so-called "characteristic factors" for ergodic averages of the form

$$\lim_{N\to\infty} \frac{1}{N^d} \sum_{n_1,\dots,n_d=1}^{N} \int T^{L_1(\mathbf{n})} f_1(x)\; T^{L_2(\mathbf{n})} f_2(x)\; \dots\; T^{L_m(\mathbf{n})} f_m(x) d\mu(x),$$

where T is a measure-preserving map on a probability measure space $(X, \mathcal{B}, \mu)$ and the functions f_i belong to $L^\infty(\mu)$. A characteristic factor is a system onto which one can project without losing any quantitative information about the average under consideration. The aim is to find characteristic factors which possess enough structure to allow one to establish convergence of the above average in a rather explicit way. For example, it was shown by Host and Kra [HK05] and Ziegler [Zie07] independently that when the linear forms $L_1, ..., L_m$ describe an arithmetic progression of length m, then there exists a characteristic

factor for the average which is isomorphic to an inverse limit of a sequence of $(m-2)$-step nilsystems. For $m = 4$, these very structured objects are closely related to the quadratic factor we are using in this paper, on which computations can be performed rather straightforwardly. After these remarks it should not be surprising that there is a notion of degree associated with a characteristic factor. What we have called the true complexity of a linear system is closely analogous to the degree of the minimal characteristic factor.

Leibman [Lei07] characterizes the degree of the minimal characteristic factor for general linear as well as certain polynomial systems. As his definition of complexity in the ergodic context is highly technical, we shall simply illustrate the analogy with our result by quoting two examples from Section 6 of his paper: In our terminology, both of the systems given by $(x+n+m, x+2n+4m, x+3n+9m, x+4n+16m, x+5n+25m, x+6n+36m)$ and the ever so slightly different $(x+n+m, x+2n+4m, x+3n+9m, x+4n+16m, x+5n+25m, x+6n+37m)$ have CS-complexity 2. However, the second one has true complexity 1 since its squares are independent, or, as Leibman puts it, because the six vectors $(1,1,1,1,1,1)$, $(1,c_1,c_2,\dots,c_5)$, $(1,d_1,d_2,\dots,d_5)$, $(1,c_1^2,c_2^2,\dots,c_5^2)$, $(1,d_1^2,d_2^2,\dots,d_5^2)$ and $(1,c_1d_1,c_2d_2,\dots,c_5d_5)$ span $\mathbb{R}^6$. (Here c_i, d_i are the cofficients of n, m, respectively, in the linear form $i+1$. Note that the special form of the ergodic average forces one to consider translation-invariant systems only, which leads to a formulation of square-independence that is particular to systems where one variable has coefficient 1 in all linear forms.)

In his proof of Szemerédi's Theorem, Furstenberg [Fur77] developed an important tool known as the *correspondence principle* which allowed him to deduce Szemerédi's combinatorial statement from the recurrence properties of a certain dynamical system. Our result in the $\mathbb{Z}_N$ case does not appear to follow from Leibman's result by a standard application of the correspondence principle. For an excellent introduction to ergodic theory and its connections with additive combinatorics, we refer the interested reader to [Kra06].

References

[Fur77] H. Furstenberg. Ergodic behavior of diagonal measures and a theorem of Szemerédi on arithmetic progressions. *J. Analyse Math.*, 31:204–256, 1977.

[Gow98] W.T. Gowers. A new proof of Szemerédi's theorem for progressions of length four. *GAFA*, 8(3):529–551, 1998.

[Gow01] W.T. Gowers. A new proof of Szemerédi's theorem. *GAFA*, 11:465–588, 2001.

[Gow04] W.T. Gowers. Quasirandomness, counting and regularity for 3-uniform hypergraphs. To appear. Available at http://www.dpmms.cam.ac.uk/∼wtg10, 2004.

[Gow06] W.T. Gowers. Two examples in additive combinatorics. Unpublished, 2006.

[Gre05a] B.J. Green. Finite field models in additive combinatorics. In *Surveys in combinatorics 2005*, volume 327 of *London Math. Soc. Lecture Note Ser.*, pages 1–27. Cambridge Univ. Press, Cambridge, 2005.

[Gre05b] B.J. Green. Montreal lecture notes on quadratic Fourier analysis. Available at arXiv:math.CA/0604089, 2005.

[GT05a] B.J. Green and T. Tao. An inverse theorem for the Gowers U^3-norm. To appear. Available at arXiv:math.NT/0503014, 2005.

[GT05b] B.J. Green and T. Tao. New bounds for Szemerédi's theorem, I: Progressions of length 4 in finite field geometries. Submitted. Available at arXiv:math.CO/0509560, 2005.

[GT06] B.J. Green and T. Tao. Linear equations in primes. Submitted. Available at arXiv:math.NT/0606088, 2006.

[HK05] B. Host and B. Kra. Nonconventional ergodic averages and nilmanifolds. *Annals of Math.*, 161(1):397–488, 2005.

[Kra06] B. Kra. Ergodic methods in additive combinatorics. Available at arXiv:math.DS/0608105, 2006.

[Lei07] A. Leibman. Orbit of the diagonal of a power of a nilmanifold. Available at http://www.math.ohio-state.edu/∼Leibman/preprints, 2007.

[Mes95] R. Meshulam. On subsets of finite abelian groups with no 3-term arithmetic progressions. *J. Combin. Theory Ser. A*, 71(1):168–172, 1995.

[Zie07] T. Ziegler. Universal characteristic factors and Furstenberg averages. *J. Amer. Math. Soc.*, 20:53–97, 2007.

Department of Pure Mathematics and Mathematical Statistics, Wilberforce Road, Cambridge CB3 0WB, UK.

E-mail address: W.T.Gowers@dpmms.cam.ac.uk

Department of Pure Mathematics and Mathematical Statistics, Wilberforce Road, Cambridge CB3 0WB, UK.

E-mail address: J.Wolf@dpmms.cam.ac.uk